GPU Developments for Applications in Climate and Weather

Stan Posey, Program Manager, Earth System Science, NVIDIA HQ, Santa Clara, CA, USA

AGENDA

- NEW NVIDIA GPU: A100
- HPC APPLICATIONS
- ML APPLICATIONS

World-Leading HPC Systems Deploy NVIDIA GPUs

ORNL Summit
#1 Top 500
27,648 GPUs| 144 PF

LLNL Sierra
#2 Top 500
17,280 GPUs| 95 PF

Piz Daint
Europe's Fastest
5,704 GPUs| 21 PF

ABCI
Japan's Fastest
4,352 GPUs| 20 PF

ENI HPC4
Fastest Industrial
3,200 GPUs| 12 PF

NERSC-9 HPC System Based Perlmutter Based on A100 GPU

Perlmutter: A System Optimized for Science

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

SC18 Gordon Bell Award: NERSC and NVIDIA Team

Exascale Deep Learning for Climate Analytics

Thorsten Kurth* tkurth@lbl.gov
Sean Treichler† sean@nvidia.com
Joshua Romero† joshr@nvidia.com
Mayur Mudigonda* mudigonda@berkeley.edu
Nathan Luehr† nluehr@nvidia.com
Everett Phillips† ephillips@nvidia.com
Ankur Mahesh* amahesh@lbl.gov
Michael Matheson‡ mathesonma@ornl.gov
Jack Deslippe* jrdeslippe@lbl.gov
Massimiliano Fatica† mfatica@nvidia.com
Prabhat* prabhat@lbl.gov
Michael Houston† mhouston@nvidia.com

Dr. Kurth now NVIDIA

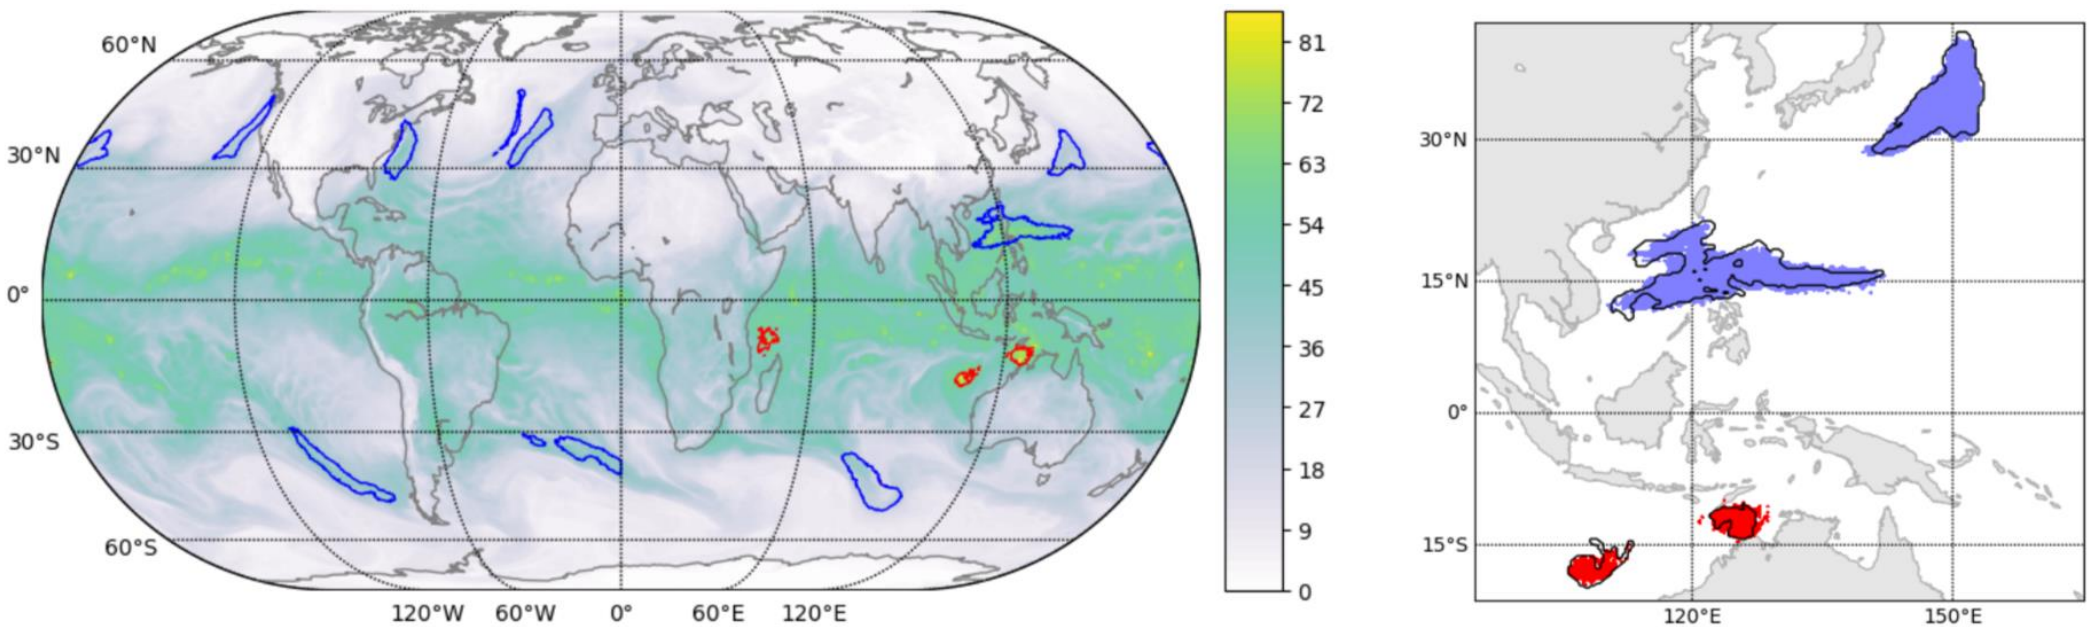

Segmentation of Tropical Storms and Atmospheric Rivers on Summit using convolutional neural networks.

Feature Progression in NVIDIA GPU Architectures

	A100 (2020)		V100 (2017)		P100 (2016)
Peak FP64 TF/s	9.7	1.3x	7.5	1.0x	5.3
Peak FP64 TC TF/s	19.5				
Peak FP32 TFlop/s	19.5		15.0		10.6
Peak TF32 TC TF/s	156				
Peak FP16 TFlop/s	312	1.7x	120	1.0x	21.2
Memory BW (GB/s)	1555	1.7x	900	1.0x	720
Memory Capacity	40 GB		16 or 32 GB		16 GB
Interconnect	NVLink: Up to 600 GB/s PCIe: 64 GB/s		NVLink: Up to 300 GB/s PCIe: 32 GB/s		NVLink: 160 GB/s PCIe: 32 GB/s
Max Power	400W		250W - 300W		300W

A100 GPU Speedups for ICON >10x vs. CPU

Single GPU Speedups vs. 2 x BDW CPU (Single Node)

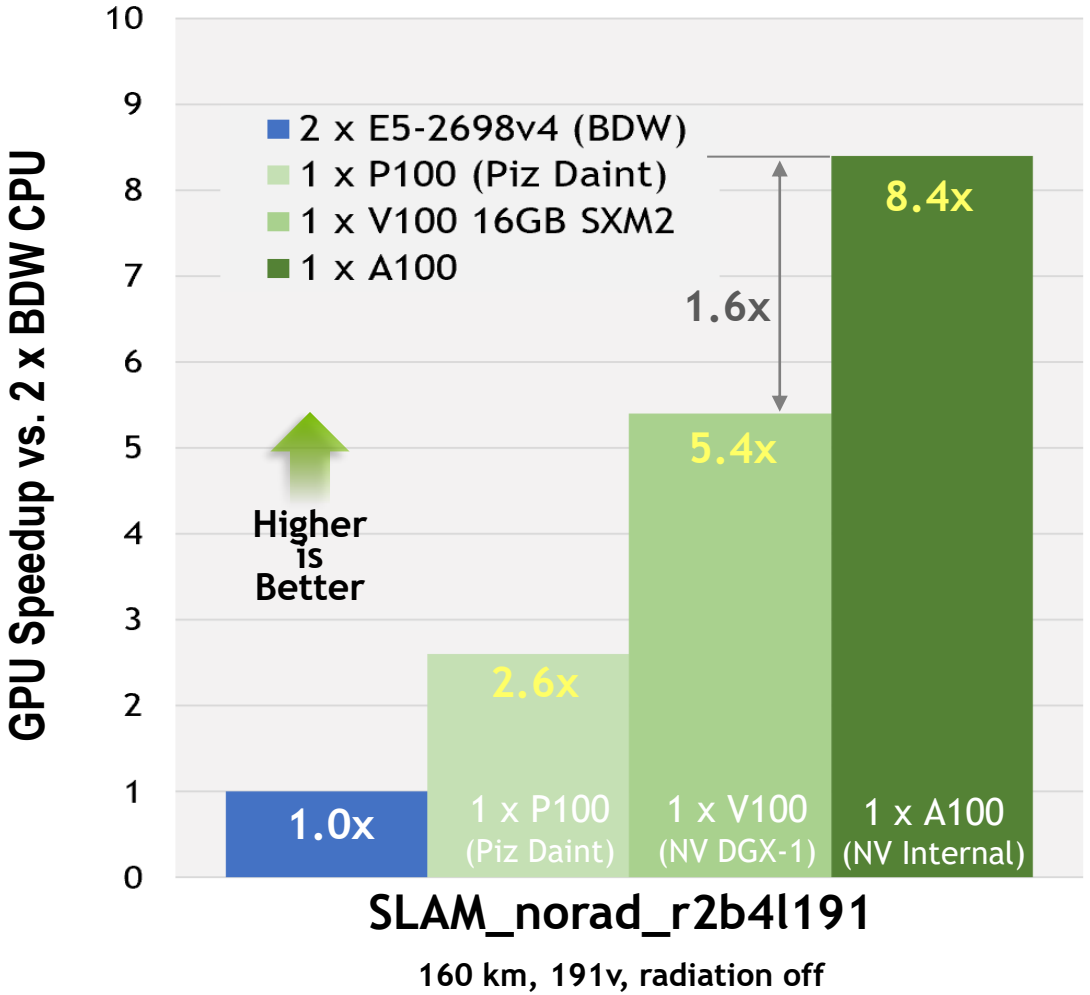

ICON on GPUs

Source: NVIDIA, May 2020
Dr. Dmitry Alexeev

Collaboration: MPI-M, DRKZ and CSCS/MCH

- End-to-end GPU with minimal data transfer
- OpenACC except use of CUB DeviceScan in CUDA to build index lists on the GPU
- Results with no land model, no radiation
- 180 model time steps

AGENDA

- NEW NVIDIA GPU: A100
- **HPC APPLICATIONS**
- ML APPLICATIONS

NVIDIA HPC Collaborations With Atmospheric Models

	Model	Organizations	Funding Source
Global:			
	E3SM-HOMEXX	US DOE: ORNL, SNL	E3SM, ECP
	MPAS-A	NCAR, UWyo, KISTI, IBM	WACA II
	FV3	NOAA	SENA
	NUMA/NEPTUNE	US Naval Res Lab, NPS	ONR
	IFS	ECMWF	ESCAPE
	GungHo/LFRic	MetOffice, STFC	PSyclone
	ICON	DWD, MPI-M, CSCS, MCH	PASC ENIAC
	KIM	KIAPS	KMA
	CLIMA	CLiMA (NASA JPL, MIT, NPS)	Private, US NSF
	FV3	Vulcan, UW/Bretherton	Private
Regional:			
	COSMO	MCH, CSCS, DWD	PASC GridTools
	AceCAST-WRF	TempoQuest	Venture backed

2019 ORNL Hackathons and GPU Model Progress

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

ABOUT OLCF ▾ OLCF RESOURCES ▾ R&D ACTIVITIES ▾ SCIENCE AT OLCF ▾ FOR USERS ▾ OLCF MEDIA ▾

Location - Date	Organizations	Model(s)	Hackathon Project
KISTI (KR) - Feb	KISTI	MPAS	Physics (WSM6)
CAS (CN) - May	CMA	GRAPES	PRM advection
ETH Zurich- Jun	MCH, MPI-M, CSCS	ICON	Physics, radiation
MIT - Jun	MIT, CliMA	CliMA Ocean	Subgrid scale LES
Princeton - Jun	NOAA GFDL	FV3GFS	SWE mini-app kernels, GFS radiation package
NERSC - Jul	DOE LBNL	E3SM	MMF (ECP)
Sheffield - Aug	Met Office, NCAS	UM, NEMO	Dycore, Miniapp
Met Office - Sep	Met Office, STFC	NEMOVAR, WW III	Miniapp (?)
ORNL - Oct	DOE ORNL, SNL	E3SM	SCREAM (Kokkos)

https://www.olcf.ornl.gov/for-users/training/gpu-hackathons/

GPU Model Progress: MPAS

- **GTC Keynote: Toward an Exascale Earth System Model with Machine Learning Components: An Update**
  Richard Loft, Director of Technology Development, CISL, NCAR **[S21834 – Recording]**

- **Articles: IBM/The Weather Company GRAF/MPAS deployment on GPUs**
  - IBM Case Study – Nov 19: IBM GRAF Weather Forecasting of the Future
  - HPC Wire – Jan 20: IBM's New Global Weather Forecasting System Runs on GPUs
  - Blog: New IBM Supercomputer Optimized for NVIDIA GPUs to Bring Better Weather Predictions Worldwide

- **Talks: AMS 2020 (Boston) presentations on MPAS GPU developments**
  - NWP at The Weather Company: Overview of a Global Rapidly-Updating Forecast System [PDF (12MB)]
    Todd Hutchinson, The Weather Company, an IBM Business, Andover, MA, et al.
  - An Implementation of MPAS-Atmosphere Running on GPUs [PDF (1MB)]
    Raghu Raj Prasanna Kumar, NVIDIA, Santa Clara, CA; and M. Duda, S. Suresh, T. Hutchinson, and J. Wong
  - Weather Forecasting - What Have We Learned and Where We Are Headed [Recording]
    Pat Feldhausen, The Weather Company, an IBM Business, Andover, MA
  - Reducing the Carbon Footprint of MPAS Weather Modeling with GPUs [PDF (2MB)]
    Rich Loft, NCAR, CISL (Presented in the NCAR Exhibit at AMS 2020)

MPAS GPU Scalability on ORNL Summit

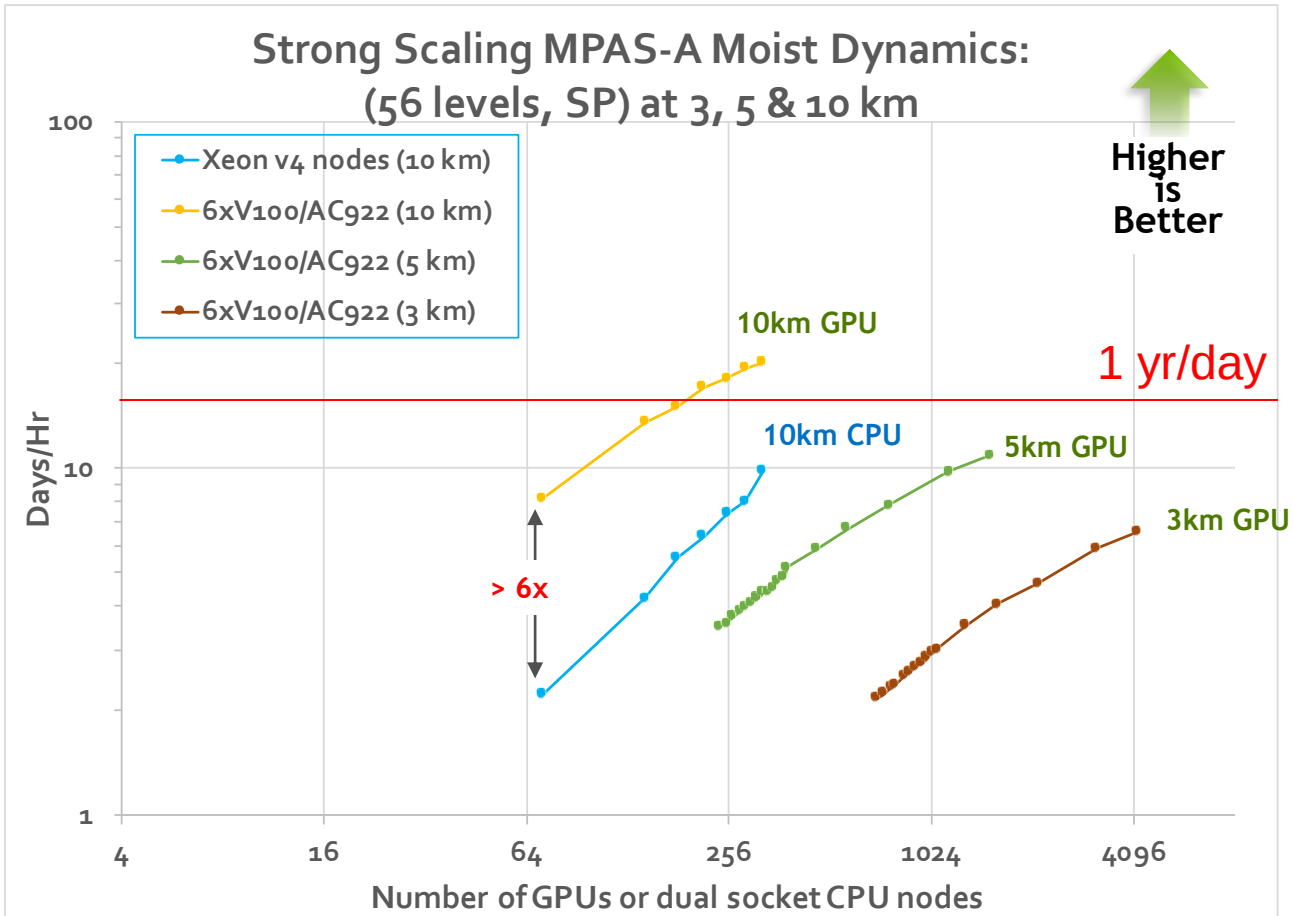

AMS 2020

12 – 16 Jan 2020, Boston, USA

An Implementation of MPAS-Atmosphere Running on GPUs

Dr. Raghu Kumar, et al.

- ORNL Summit GPU system (V100)
- NCAR Cheyenne CPU system (BDW)

An Implementation of MPAS-Atmosphere Running on GPUs [PDF (1MB)]

IBM/TWC GRAF Model Based on GPU MPAS

IBM Makes Higher Quality Weather Forecasts Available Worldwide

ARMONK, N.Y. and ATLANTA, Nov. 14, 2019 /PRNewswire/ -

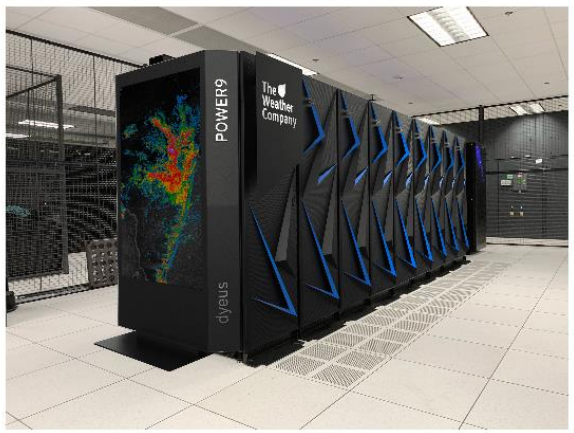

TWC operational configuration of 3km – 15km with hourly updates using ~300 x V100 GPUs

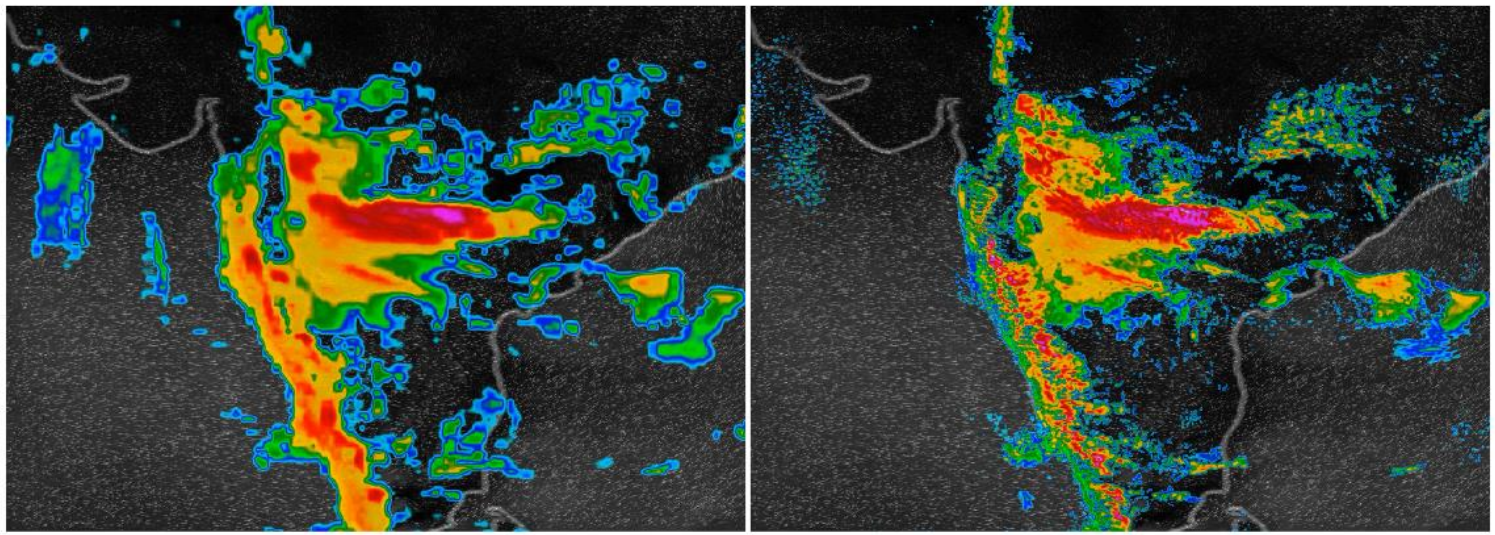

https://www.hpcwire.com/2019/01/09/ibm-global-weather-forecasting-system-gpus/

MeteoSwiss GPU System for NWP to use ICON

FRANKFURT, Germany, June 18, 2019 (GLOBE NEWSWIRE) -

Swiss Federal Office of Meteorology and Climatology Advances Weather Forecasting With New Cray Supercomputer and Storage

CSCS' new CS-Storm® is configured with 18 compute nodes, each with 8 NVIDIA® V100 GPUs and 2 Intel® Xeon® Gold 6134 CPUs, and includes two Cray ClusterStor® L300 storage systems. The CS-Storm was accepted in April 2019 and will become fully operational in 2020.

18 Nodes x 8 x V100 = 144 Total V100 GPUs

MeteoSwiss Roadmap

- V100 system since 2019
- New EPS configurations operational during 2020
- New ICON-LAM in ~2022 (Pre-operational 2020)

MeteoSwiss COSMO NWP Configurations During 2020

With V100 GPUs

IFS from ECMWF

4 per day, 18km / 9km (?)

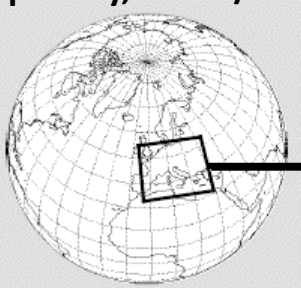

COSMO-2E (2 KM)

4 per day, 5 day forecast

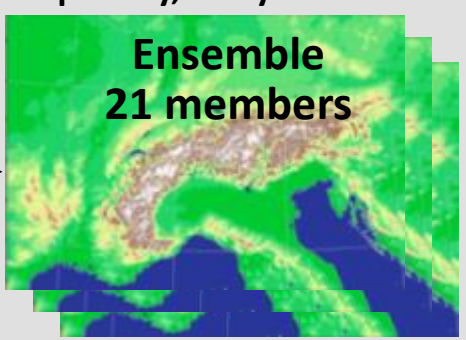

COSMO-1E (1 KM)

8 per day, 33 hr forecast

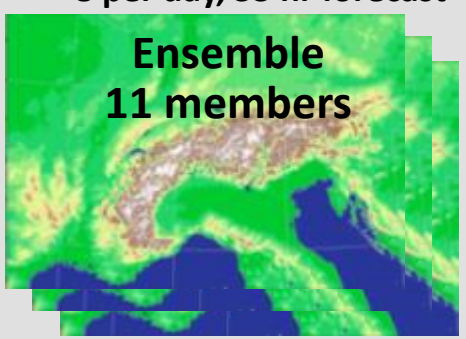

GPU Model Progress: IFS

ECMWF

- **DOE INCITE 2020**: ECMWF scientists to simulate global weather at 1 km resolution
  - *The project will make it possible to simulate deep convection explicitly at the global scale.*
  - US DOE allocation of 500,000 node-hours on the Summit supercomputer at ORNL.
  - GPU evaluations at-scale of developments from the ECMWF Scalability Programme
- **Goal: Explore the limits of IFS-spectral vs. IFS-FV models at large scale**
  - Initial experiments to test largest CPU allocations possible (Summit has 9,200 x P9 CPUs).
  - Next experiments with model components on the GPU: spectral transforms, physics, advection.
  - Blog – Jun 19: Optimising the ECMWF Integrated Forecasting System (IFS) on the Summit computer
- **Other weather and climate projects using the Summit GPU system**
  - CPU-based 25 km climate model combined with GPU “super-parametrised” cloud models
  - Exascale deep learning for climate analytics: image classification, localisation, object detection
  - Experiments with Accelerating High-Resolution Weather Models with Deep-Learning Hardware
    - PASC Conference 2019 Best paper award; Recording ; Paper ; Paper download

ECMWF IFS Dwarf Optimizations – Single-GPU

Hybrid Computing – single GPU

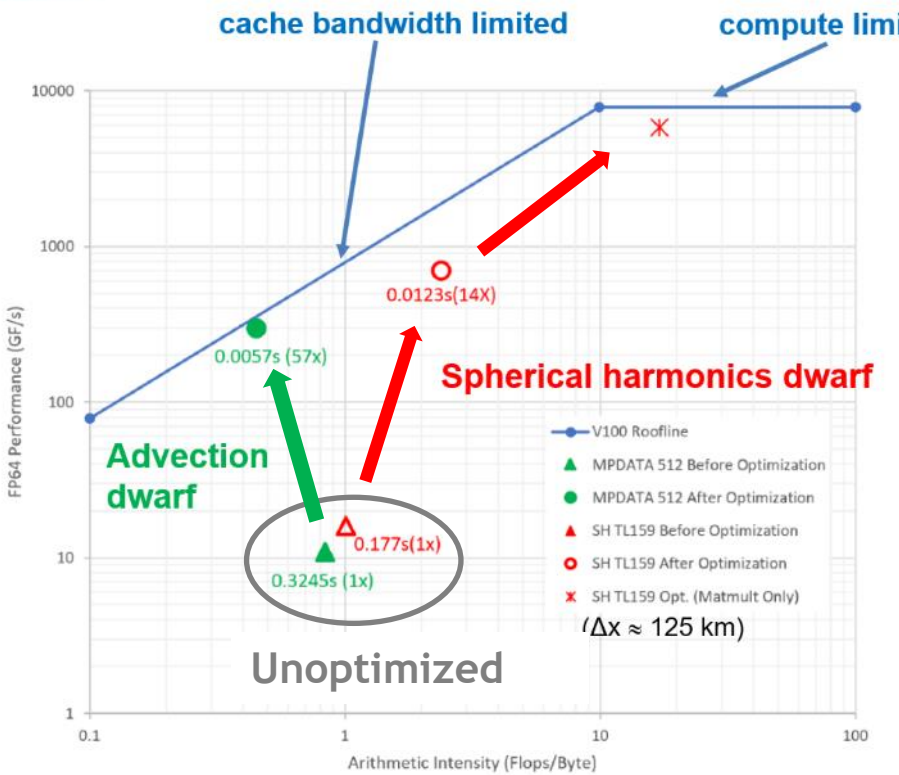

SH Dwarf = 14x

Advection Dwarf = 57x

by:

- exposing parallelism in loops for OpenACC mapping
- Kernel optimization by memory mapping
- exploiting CUDA BLAS features
- minimizing data allocation and movement

From *"ECMWF Scalability Programme"*

Dr. Peter Bauer, UM User Workshop, MetOffice, Exeter, UK 15 June 2018

- **Single V100 GPU improved SH dwarf by 14x vs. original**
- **Single V100 GPU improved MPDATA dwarf 57x vs. orig**

ECMWF IFS SH Dwarf Optimization – Multi-GPU

Hybrid Computing – multiple GPU

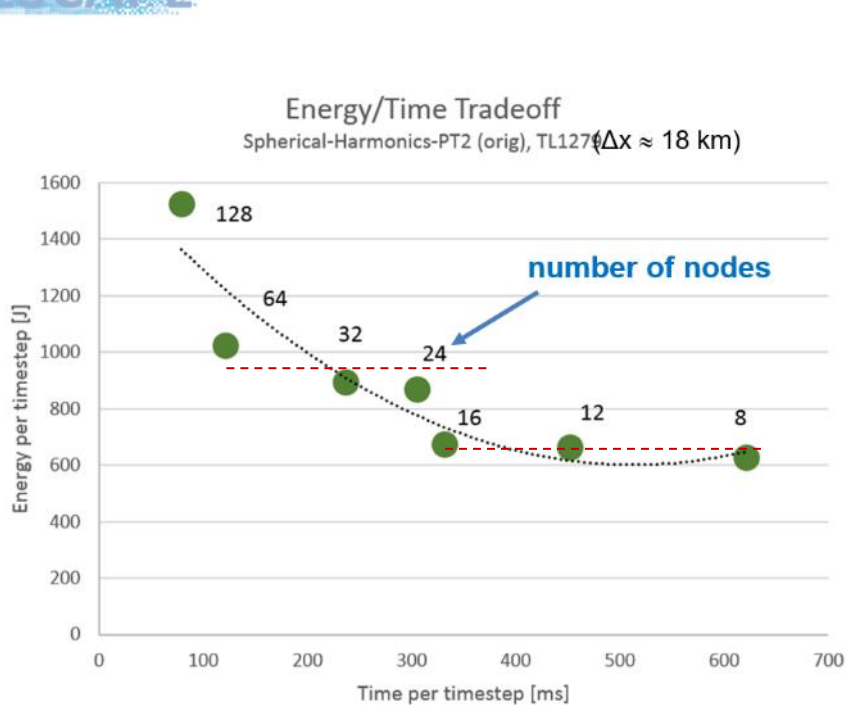

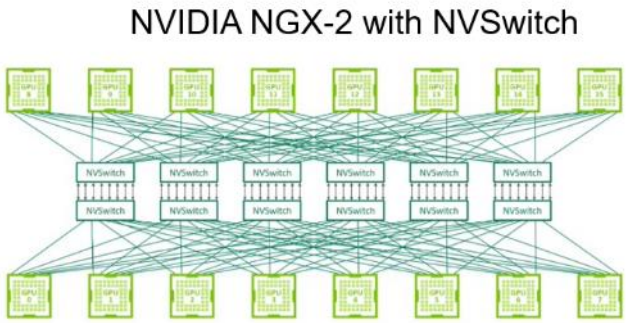

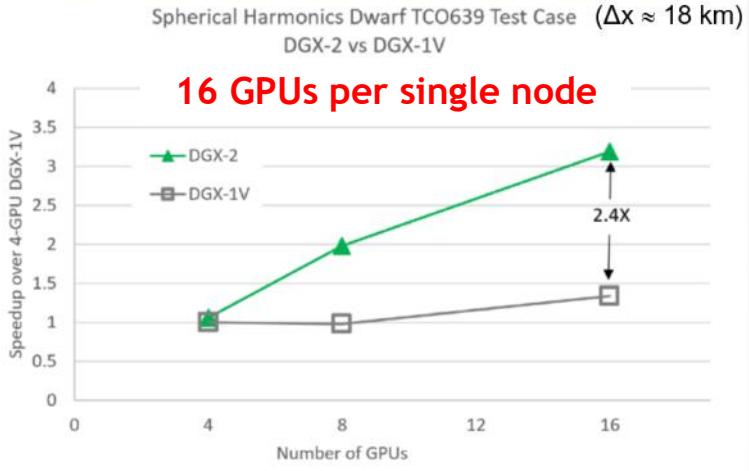

From *"ECMWF Scalability Programme"*

Dr. Peter Bauer, UM User Workshop, MetOffice, Exeter, UK 15 June 2018

- Results of Spherical Harmonics Dwarf on NVIDIA DGX System
- Additional 2.4x gain from DGX-2 NVSwitch for 16 GPU systems

IFS Scaling to 11,520 x GPUs on ORNL Summit

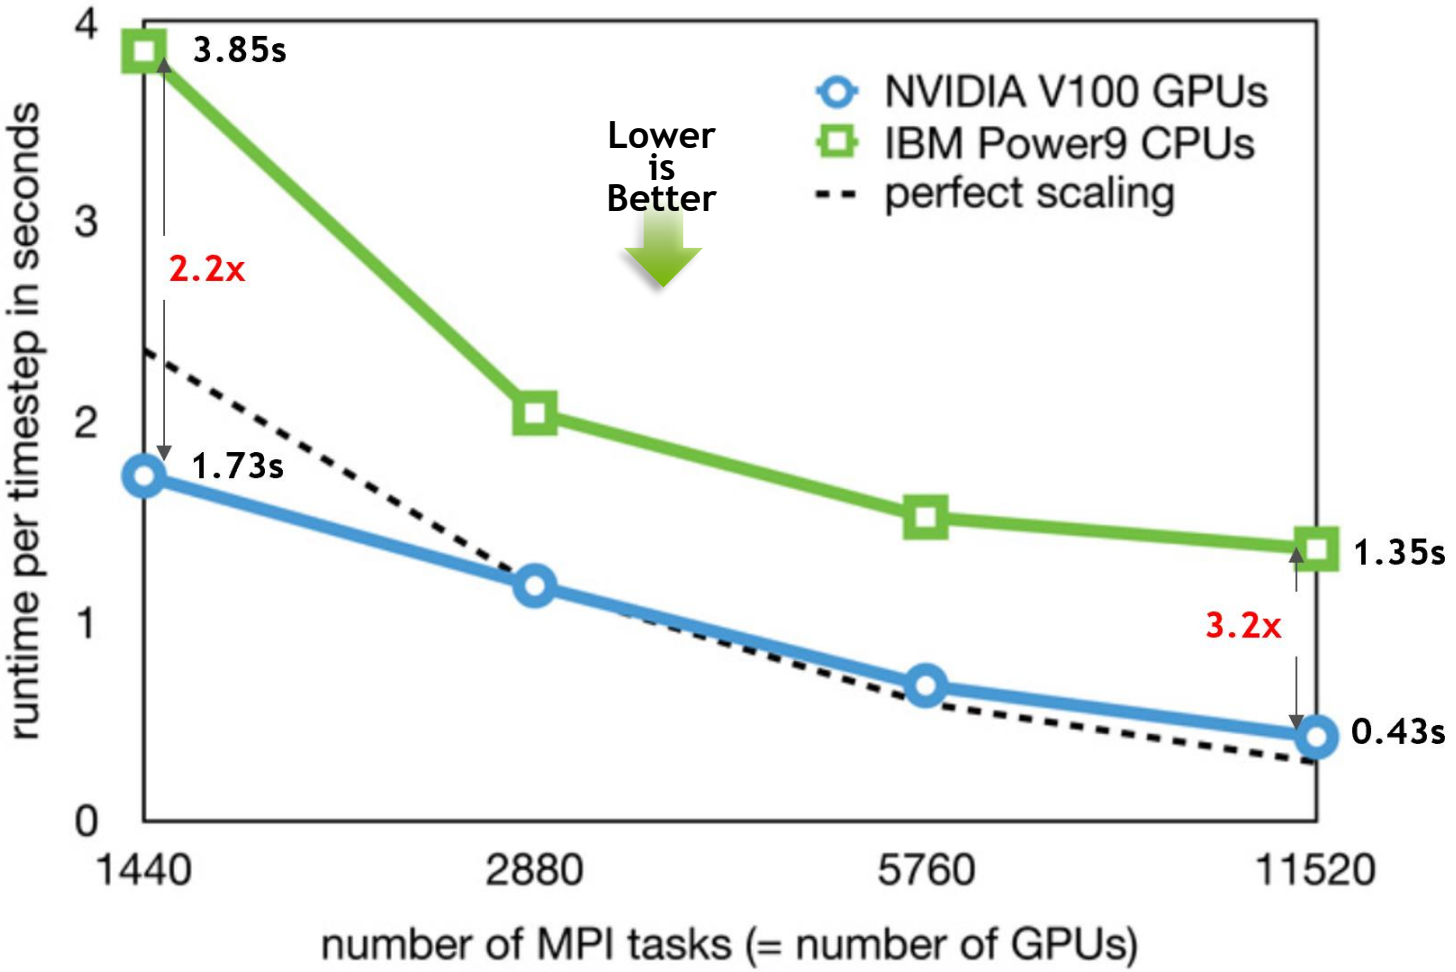

ECMWF

Scaling of spherical harmonics dwarf across multiple Summit nodes, using a hybrid OpenMP/OpenACC/MPI configuration, using GPUdirect and CudaDGEMM/FFT libraries. 11520 MPI tasks in this figure use 1920 nodes.

https://www.ecmwf.int/en/about/media-centre/science-blog/2019/oak-ridge-leadership-computing-facility-user-workshop-2019

ECMWF Study of Half Precision on GPUs

ECMWF

Can we use deep learning hardware for conventional models?

Relative cost for model components for a non-hydrostatic model at 1.45 km resolution:

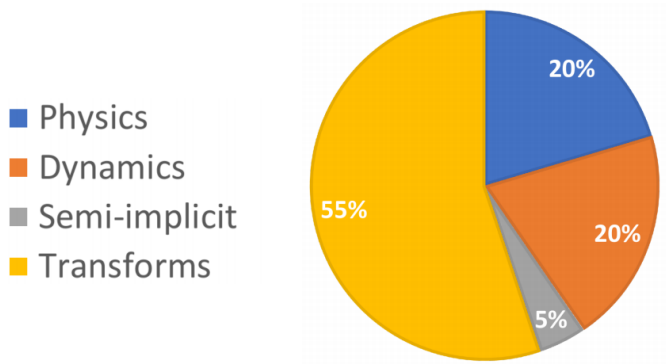

- The Legendre transform is the most expensive kernel. It consists of a large number of standard matrix-matrix multiplications.
- If we can re-scale the input and output fields, we can use half precision arithmetic.
- Tensor Cores on NVIDIA Volta GPUs are optimised for half-precision matrix-matrix calculations with single precision output. 7.8 TFlops for double precision vs. 125 TFlops for half precision on the Tensor Core.

Half precision Legendre Transformations

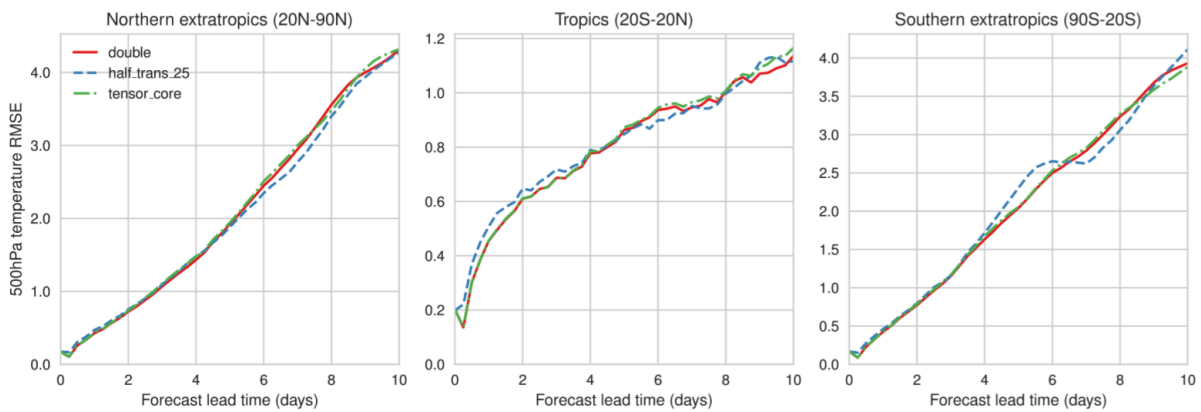

Root-mean-square error for geopotential height at 500 hPa at 9 km resolution averaged over multiple start dates.

- The simulations are using an emulator to reduce precision (*Dawson and Dueben GMD 2017).*
- More thorough diagnostics are needed.

Hatfield, Chantry, Dueben, Palmer Best Paper Award PASC2019

Machine Learning for Weather Forecasts

Dr. Peter Dueben, ECMWF
Keynote Talk
2020 Stanford Conference

https://www.hpcadvisorycouncil.com/events/2020/stanford-workshop/pdf/PeterDueben_MachineLearnWeatherForecast_StanfConf20_042120.pdf

NVIDIA

AGENDA

- NEW NVIDIA GPU: A100
- HPC APPLICATIONS
- **ML APPLICATIONS**

Select ML Applications in Weather and Climate

DETECTION

- **Tropical storms**

Extra-tropical cyclones
Atmospheric rivers
Cyclogenesis events
Convection initiation
Change detection

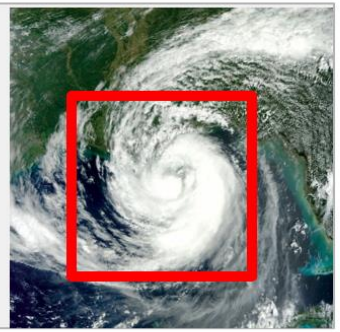

ENHANCEMENT

Frame repair
Sequence repair
- **Slow motion**

Super-resolution
Cloud removal
Data augmentation

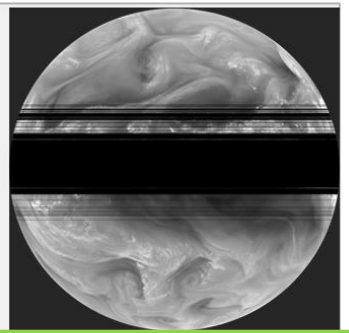

TRANSLATION

- **Data Assimilation**

Forecast verification
Model inter-comparison
Common data formatting
Colorization
Digital Elevation from Imagery

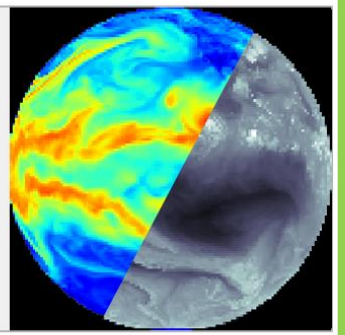

EMULATION

Physics Acceleration
Turbulence
- **Radiation**

Convection
Microphysics
Dynamics Acceleration

PREDICTION

Uncertainty prediction
Storm track
Storm intensity
Fluid motion
Now casting
Satellite frame prediction

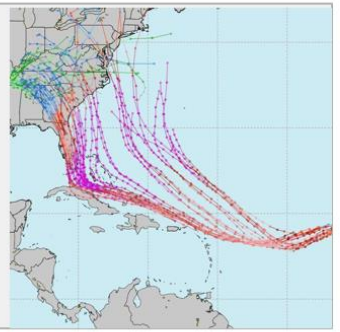

PARAMETRIZATION

New parametrizations
From higher resolution model
- **From observational data**

ECMWF Collaboration on ML Emulation of Radiation

Machine Learning for Weather Forecasts

ECMWF

Numerical weather forecasts: To emulate the radiation scheme

- Store input/output data pairs of the radiation schemes
- Use this data to train a neural network
- Replace the radiation scheme by the neural network within the model

This is a very active area of research:
Rasp, Pritchard, Gentine PNAS 2018
Brenowitz and Bretherton GRL 2018
…

Why would you do this?
Neural networks are likely to be much more efficient and portable to heterogenous hardware

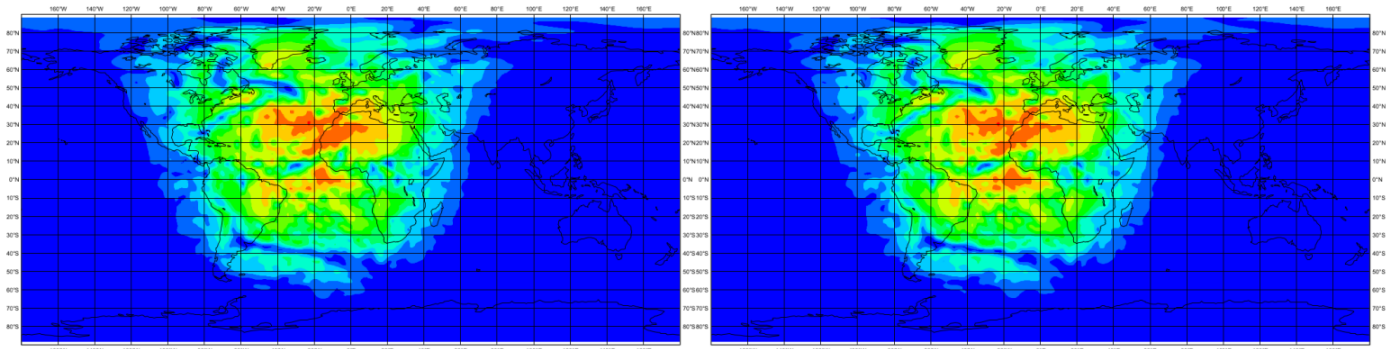

Surface downward solar radiation for the original scheme and the neural network emulator (based on a ResNet).

The approach is working and the neural network is ~10 times faster than the original scheme. However, model results are still degraded.

Dueben, Hogan, Bauer @ECMWF and Progsch, Angerer @NVIDIA

From *"Machine Learning for Weather Forecasts"*

Dr. Peter Dueben
2020 Stanford Conference, Apr 2020

- ResNet34, Resnet18 networks on-line IFS show agreement with reference model to multi day forecasts, with good speedups
- Next will experiment with grid resolution and frequency of radiation calls.

https://www.hpcadvisorycouncil.com/events/2020/stanford-workshop/pdf/PeterDueben_MachineLearnWeatherForecast_StanfConf20_042120.pdf

NVIDIA

DOE Collaboration on ML Emulation of Radiation

ANN Emulation of RRTMG Provides 10x Speedup

-From AGU 2017 by M. Norman and P. Anikesh, Oak Ridge NL

Rapid Radiative Transfer Model for GCMs

E3SM Energy Exascale Earth System Model

Surface Net SW Flux (RRTMG). Mean = 161.91 W/m^2

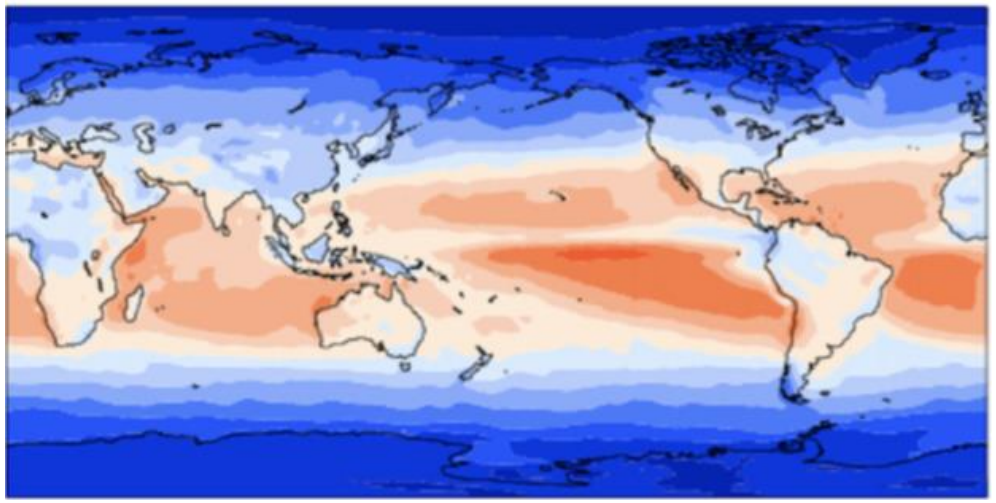

Surface Net SW Flux (Emulation). Mean = 161.91 W/m^2

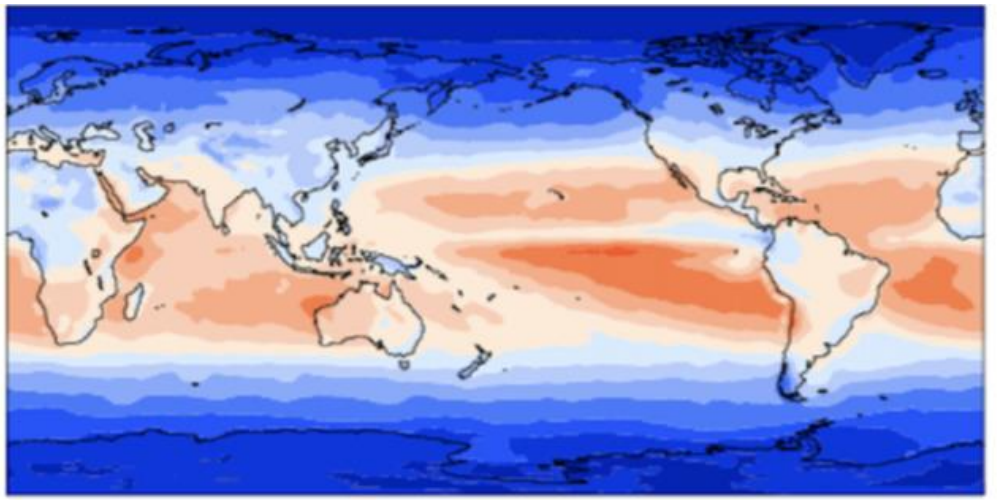

Emulation of radiative transfer parametrization
E3SM global climate model
Speedup of 8-10x over original.
Details: 3778 inputs, fully connected, 3 hidden layers, 6million training samples

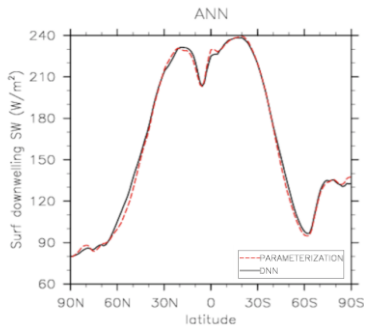

NASA Collaboration on ML-Based AQM Research

Atmospheric Chemistry Modeling and Air Quality Forecasting using Machine Learning

Christoph A. Keller
NASA Global Modeling and Assimilation Office (GMAO)
Universities Space Research Association (USRA)

Mat J. Evans
Wolfson Atmospheric Chemistry Laboratories, University of York
National Centre for Atmospheric Sciences, University of York

1st NOAA Workshop on Leveraging AI
23-25 April 2019

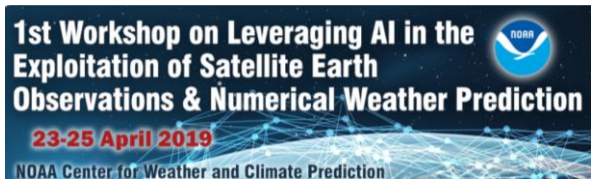

NASA USRA UNIVERSITY of York National Centre for Atmospheric Science NATURAL ENVIRONMENT RESEARCH COUNCIL

NVIDIA

NASA Collaboration on ML-Based AQM Research

Atmospheric chemistry models are computationally very expensive

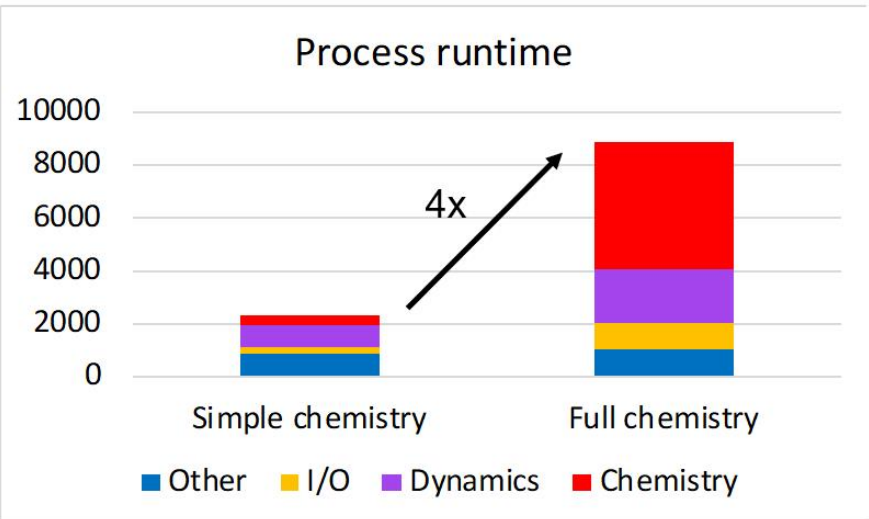

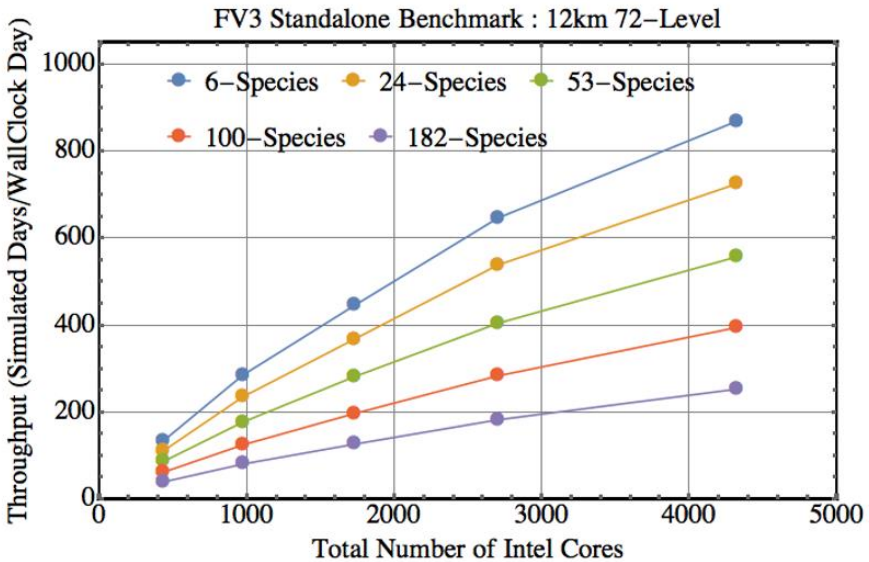

Courtesy of W. Putman, NASA GMAO

- High-resolution chemistry simulation requires >1000 CPU's
- Can simulate approx. 20 days in 24 hours
- Outputting the full chemical state is 1.5 TB / simulation day

NASA Collaboration on ML-Based AQM Research

Replace chemical integrator with machine learning model

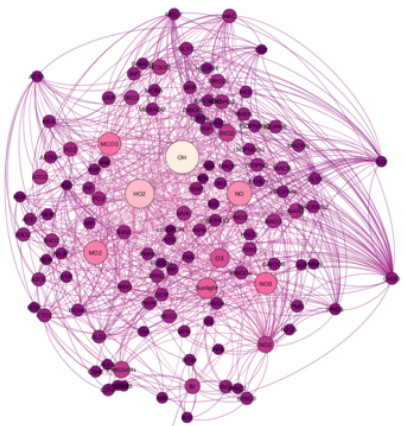

Numerical model

Initialize
Transport
Emissions
Sub-grid
Photolysis
Chemical Integrator
Deposition
Diagnostics

Training data

Random forest

Emulator

Initialize
Transport
Emissions
Sub-grid
Photolysis
Deposition
Diagnostics

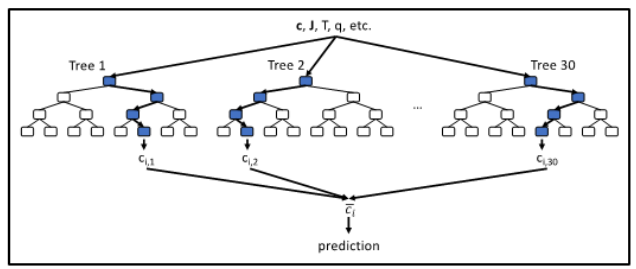

NASA Collaboration on ML-Based AQM Research

Random forest / XGBoost training benchmarks

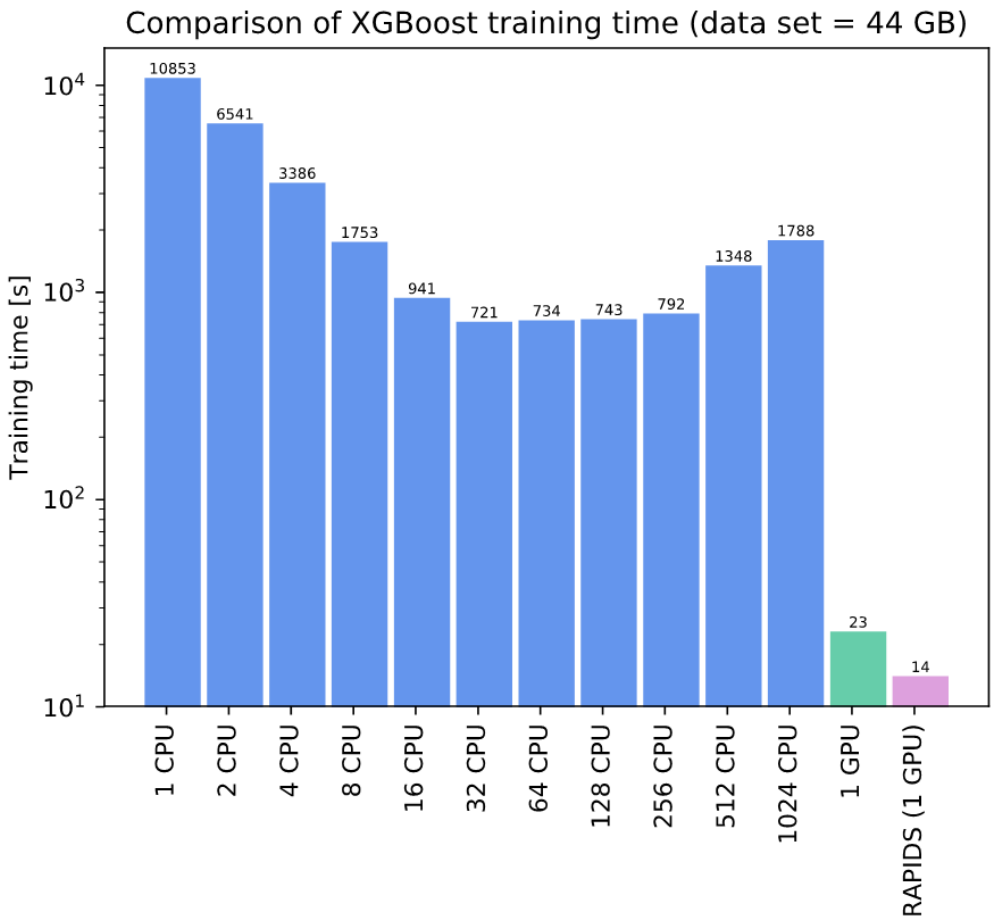

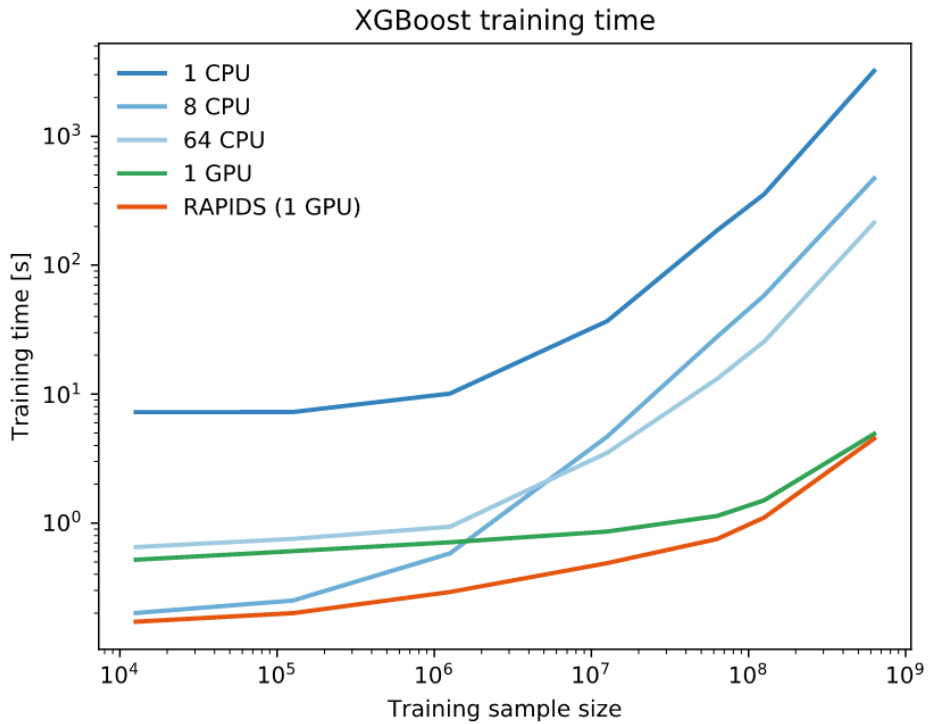

NASA Collaboration on ML-Based AQM Research

Random forest overestimates ozone surface concentrations over remote regions

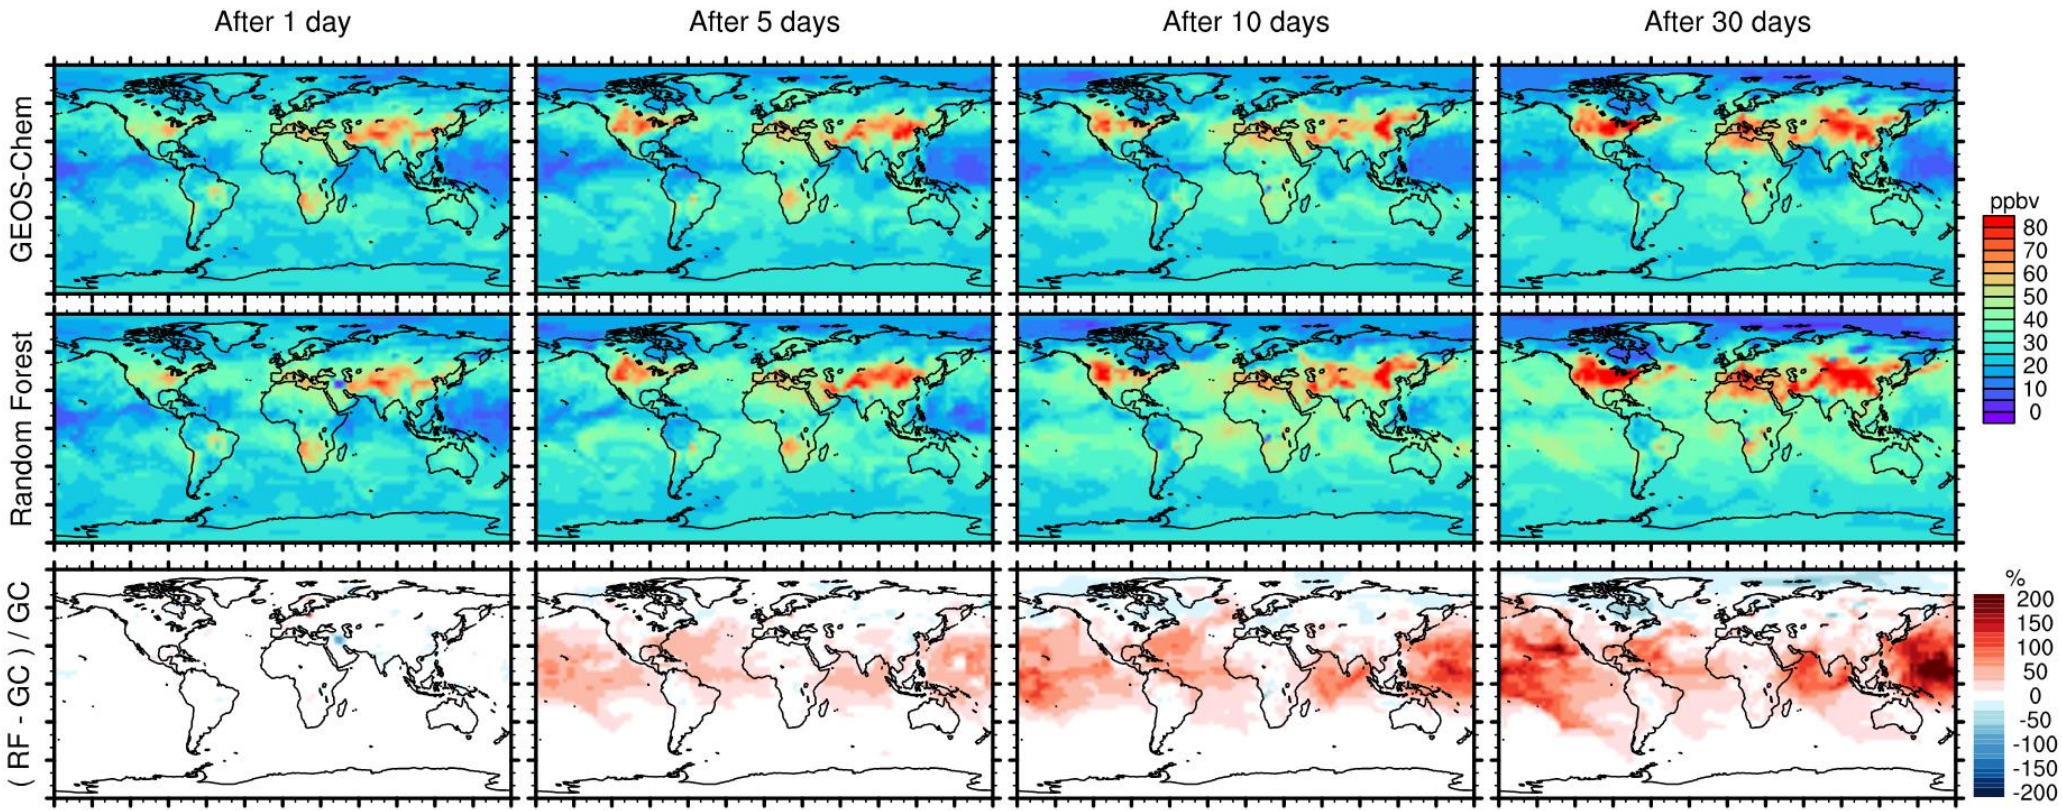

SUMMARY

- OPPORTUNITIES FOR APPLICATION BENEFITS FROM NEW A100 GPU
- GOOD PROGRESS IN NUMERICAL MODEL DEVELOPMENT ON GPUS
- FAST GROWTH IN ML APPLICATIONS OWING TO AI FEATURES OF GPU

Thank you and Questions?

Stan Posey, sposey@nvidia.com